ELSEVIER

Behavioural Processes 44 (1998) 163–182

BEHAVIOURAL PROCESSES

Are separate theories of conditioning and timing necessary?

Kimberly Kirkpatrick *, Russell M. Church

Department of Psychology, Brown University, Providence, RI 02912, USA

Received 30 March 1998; received in revised form 18 August 1998; accepted 28 August 1998

Abstract

Conditioning and timing studies have evolved under separate traditions, which is exemplified in both traditional theories (e.g. the Rescorla–Wagner model of conditioning vs. Scalar Timing Theory) and in a dual process model (Gibbon, J., Balsam, P., 1981. In: Autoshaping and Conditioning Theory. Academic Press, New York.). Other lines of theoretical development in both timing and conditioning fields have resulted in the emergence of 'hybrid' theories in which conditioning and timing processes are integrated. Simulations were conducted with a recent hybrid theory of timing (Machado, A., 1997. Psychol. Rev. 104, 241–265). The simulations were of classical conditioning procedures in which the local or global predictability of food was varied by manipulating the variability of the CS–US relationship, variability of the CS duration, and variability of the intertrial interval. The hybrid model provided good qualitative fits to indices of conditioning (discrimination ratios) and timing (local rates of responding), indicating that it may be possible to model both conditioning and timing results with a single process in which an internal representation of time and a strength of association are integrated. However, the failure of the model to provide good quantitative fits of the data indicates the need for a consideration of alternative perceptual representations of time and/or principles of association within the framework of the hybrid model.

Keywords: Conditioning; Rats; Theory; Timing

1. Introduction

Conditioning and timing studies have developed quite separately. They differ in the independent variables investigated, in the intervening variables that are assumed, and in the dependent variables that are measured. For example, in conditioning, the CS–US contingency may be used as an independent variable; in timing, stimulus duration may be used as an independent variable. In conditioning, associative strength may be used as an intervening variable; in timing, remembered duration may be used as an intervening variable.

* Corresponding author. Tel.: +1 401 8633979; fax: +1 401 8631300; e-mail: kim_kirkpatrick@brown.edu

0376-6357/98/$ - see front matter
PII: S0376-6357(98)00047-3

In conditioning, CR magnitude may be used as a dependent variable; in timing, the times of responses may be used as a dependent variable. Although these differences appear to be substantial, it is possible that none of the apparent differences are fundamental.

A standard timing experiment may be described in the following way: a signal (e.g. keylight) is turned on; at some fixed duration after signal onset (e.g. 15 s), food is made available; the first response following the time of food availability results in both food delivery and signal termination; an intertrial interval of a long and variable duration intervenes between the previous food presentation and the next signal onset. This series of events, known as the discrete-trials fixed interval procedure, normally results in an increasing rate of responding as a function of time since signal onset, with a maximum response rate occurring near the time when food is available.

A standard conditioning experiment may be described in the following way: a signal (e.g. keylight) is turned on; at some fixed duration after signal onset (e.g. 15 s), food is delivered and the signal is terminated; an intertrial interval of a long and variable duration intervenes between the previous food presentation and the next signal onset. This series of events, known as a delay conditioning procedure, normally results in growth in the strength or probability of responding during the signal as a function of some amount of training.

These two procedures differ in only one way—the operant response contingency. Specifically, in the timing procedure, food was delivered contingent on a response from the animal after a particular amount of time had passed. In the conditioning procedure, food was delivered contingent only on the amount of time that had passed since signal onset. Although the response contingency is not a trivial difference, similar aspects of the two procedures may be learned (e.g. the onset of a signal predicts that food will occur in 15 s; see, for example, Dickinson and Mackintosh, 1978; Dickinson, 1980).

In fact, evidence from both classical and instrumental procedures indicates that the general principles of response timing are similar. For example, in both cases, responses are maximally likely near the time of reinforcement; this is the 'inhibition of delay' phenomenon described by Pavlov (1927). Additionally, responses engendered by both classical (e.g. Killeen et al., 1978) and instrumental (e.g. Church et al., 1994) procedures are scalar in nature.

Given that timing and conditioning procedures are so similar and that the data generated by the two procedures share many common features, it seems plausible that the same process may be involved in generating responses. In fact, timing accounts have been used to explain autoshaped keypecking behavior evoked by classical procedures (Durlach, 1989). These theoretical explanations have incorporated the notion that conditioned responding is a function of the comparison of waiting times until food during the CS versus during the background (Gibbon and Balsam, 1981; Jenkins et al., 1981). Although temporal content is proposed to play a role in conditioning, these accounts do not produce timed responses; their output is associative (response) strength.

In the preponderance of animal learning studies, conditioning and timing still continue to receive separate treatment (see Holder and Roberts, 1985; Cole et al., 1995; Barnet and Miller, 1996; Barnet et al., 1997 for exceptions), which has resulted in the emergence of dominant theories of conditioning and timing that have little relationship to one another.

An excellent example of a conditioning theory is the Rescorla–Wagner model (Rescorla and Wagner, 1972), which states that reinforcement leads to a change in associative strength of each stimulus i, ΔV_i that is proportional to the discrepancy between the strength of the US (λ) on the present trial and the expected strength of the US (Y, the sum of the associative strength of the stimuli that are present on the current trial). That is, $\Delta V_i = \beta(\lambda - Y)\ \alpha_i X_i$, with α_i a constant for the salience of each stimulus, β a constant for the salience of the US, X_i indicating the presence ($X_i = 1$) or absence ($X_i = 0$) of stimulus i, and $Y = \Sigma(V_i \cdot X_i)$. Under conditions of positive reinforcement, λ is set to a positive value, usually 1, with nonreinforcement set to 0, and with negative

reinforcement set to a negative value. V_i can be either positive (excitatory conditioning) or negative (inhibitory conditioning). With the addition of a rule to translate strengths of association into observable response measures, it is possible to make predictions about the average rate (or probability) of responding in the presence of different stimulus configurations. Typically, however, it is only assumed that responding relates ordinally to the strength of association.

An excellent example of a timing theory is Scalar Timing Theory, which contains three interrelated modules (Gibbon et al., 1984). The first module is referred to as the clock/accumulator process, and it provides the perceptual representation of a duration. At the onset of the event that begins the timing process, a switch closes (with some variable latency), and pulses from a pacemaker reach an accumulator. The switch opens (with some variable latency) at a second time marker (such as stimulus termination or reinforcement). The sum of the pulses in the accumulator serves as the perceptual representation of time. The second module is a reference memory, and it provides the memory of previously reinforced durations. At the time that an animal receives reinforcement the number of pulses in the accumulator is transferred to reference memory, perhaps with a variable memory storage constant. Temporal memory consists of a distribution of these remembered durations—it is an examplar memory. The third module is a comparator; it provides the basis for a decision among responses (including whether or not to respond). The comparison is between a representation of current time (from the first module) and a random sample of a single remembered time (from the second module); the comparison is done with a ratio rule rather than a difference rule. With estimates of the parameters representing various sources of variance in perception, memory, and decision processes, it is possible to make predictions about when responses will occur.

The Rescorla–Wagner model and Scalar Timing Theory differ along several dimensions. First, they differ in what is perceived by the model (or animal). The Rescorla–Wagner model perceives states which identify which stimuli (contextual stimuli, CSs, and USs) are present. In contrast, Scalar Timing Theory perceives events, which are changes in the status of a stimulus (e.g. a light went on, a tone went off, food occurred); these events serve to start, reset, or stop the clock. A second difference is the learning mechanism. The Rescorla–Wagner model contains a conditioning mechanism that results in a change in the strength of association as a function of the present state. Scalar Timing Theory contains a timing mechanism that records the passage of time between events (signal to reinforcer, reinforcer to reinforcer) in reference memory. A third difference is the decision process or output of the model/animal. The Rescorla–Wagner model's output is a decision of whether or not a response will be made in the presence of a particular state. Scalar Timing Theory's output is a decision of when responses will be initiated relative to an event. Thus, in their original forms, these two models are restricted to their respective domains. The Rescorla–Wagner model cannot perceive, encode, store, or respond to temporal intervals because there is no timing mechanism. Scalar Timing Theory has no principled rules for determining which stimuli will be timed, so it cannot model stimulus selection phenomena such as blocking and overshadowing because there is no conditioning mechanism.

In the conditioning literature, a change in response strength as a function of training is the primary index of conditioning. Indices of response strength include absolute and relative response rates, probabilities of responding, the amplitude of a response, and the speed of a response. In the timing literature, response timing measures such as local rates of responding as a function of time since some event, or characterizations of the times of individual response bursts are the primary indices of timing. Although a conditioning theory and a timing theory may be used to explain the behavior produced by some of the same procedures, they are used to explain different aspects of that behavior. A conditioning theory is used to explain the average or relative response tendency in the presence of a stimulus configuration; a timing theory is used to explain the time of occurrence of the responses.

Animals produce individual responses at particular times. The response stream generated by an animal may be complex in structure, possessing many different attributes that can be assessed with many different dependent measures. For example, in research on animal timing, it has been observed that responses on a fixed interval schedule occur in a break-run pattern, with a burst of responding near the anticipated time of reinforcement (e.g. Schneider, 1969; Gibbon and Church, 1990). Many different measures may be used in developing an explanation of the response bursts, such as the overall rate of responding within a burst, the beginning and end times of the burst, the pattern of responding within a burst (e.g. steady vs. increasing rate), and the periodicity between response bursts. Each of these different dependent measures may be separately influenced by different experimental operations, or they may covary. Thus, in order to adequately understand the nature of the response stream, one needs to examine several different measures. Any theory that restricts its domain to a single response measure (e.g. the Rescorla–Wagner model only generates an overall strength in each state) will be unable to account for the complexity of behavior that is produced under various experimental manipulations.

The real-time theories of conditioning (e.g. Sutton and Barto, 1981, 1990; Blazis et al., 1986; Moore et al., 1986; Tesauro, 1986; Klopf, 1988; Moore and Desmond, 1992) are more recent extensions of the Rescorla–Wagner model that can generate timed responses as well as changes in overall response strength. The Sutton and Barto (1981) real-time model was a direct extension of the Rescorla–Wagner model that incorporated two changes: (1) the value of V_i was updated at every time step both within and between trials, as opposed to the Rescorla–Wagner model that only updated after every trial; and (2) the presence or absence of a stimulus (X_i) was replaced by a graded eligibility trace that increased during the CS and decreased after CS termination. With the addition of an eligibility trace (i.e. a clock), the Sutton and Barto real time model produced responses in real time as well as accommodating many of the effects of temporal variables in conditioning (e.g. CS–US interval effects, trace conditioning). This model and other real-time models are 'hybrid' theories because the timing and conditioning mechanisms are integrated to produce the ultimate output of the system.

Hybrid theories have emerged in the timing literature as well: two examples are the Multiple Oscillator Model (Church and Broadbent, 1990) and a recent version of Behavioral Theory of Timing (Machado, 1997). These models are extensions of earlier timing theories (Scalar Timing Theory and Behavioral Theory of Timing) that have incorporated a strength of association along with a time-keeping device. Although they have arisen to account for different types of data from different kinds of experiments, the real-time conditioning models and hybrid timing models have converged on a similar solution—an integration of timing and conditioning mechanisms.

The purpose of the present experiment was to determine whether a hybrid timing theory, Machado's (1997) variant of behavioral theory of timing, could explain the effects of variability of the stimulus-food relationship, variability of the stimulus duration, and variability of intertrial interval on both response strength and response timing measures of performance. These particular manipulations were chosen because it was expected that they would produce different effects on the conditioning (response strength) and timing (response timing) measures of performance.

One of the most basic facts of conditioning is that when a CS is followed by a US (delay conditioning), acquisition of responding occurs (Pavlov, 1927). Because most delay conditioning studies involve a fixed duration CS that is followed immediately by the US, it is possible for timing of the US to occur. So, it is possible to observe both timing and conditioning of the CS (e.g. inhibition of delay; Pavlov, 1927). On the other hand, when the CS and US presentations are uncorrelated (truly random control), no conditioning occurs (Rescorla, 1967). Moreover, because the US occurs at random times in the CS, timing of the US should not occur. So, in a truly random control, neither conditioning or timing should occur. Thus, the standard comparison between delay conditioning and random control

groups does not allow for a separate assessment of the role of conditioning and timing in the production of CRs.

In order to disentangle the influences of conditioning and timing on responding in a delay conditioning procedure, one can contrast conditions in which timing may or may not occur, without any effects on the degree of conditioning. This may be done by comparing the effects of a fixed and a variable duration CS. With a variable duration CS (in a delay conditioning paradigm), conditioning should occur to a similar degree as with a fixed duration CS (Kamin, 1960; Libby and Church, 1975). However, because a variable duration CS does not predict the time of occurrence of the US, timing should not occur during the CS (Libby and Church, 1975). Thus, the use of a fixed vs. a random duration CS allows for a comparison of conditions in which timing should (fixed duration CS) or should not occur (random duration CS), but the degree of conditioning is unaffected.

Another approach is to assess the influence of timing on responding, separate from the influence of conditioning. This may be done by contrasting the effects of fixed and random intertrial interval (ITI) durations in a delay conditioning paradigm in which the CS is a variable duration. When the ITI is fixed, then anticipation of the upcoming US can occur, with responses timed relative to the prior US (temporal conditioning; Pavlov, 1927). When the ITI is random, then no temporal conditioning should occur in the ITI. Supposedly, the degree of conditioning to the CS is unaffected by the degree of variability of the ITI (Gibbon et al., 1977). Moreover, with a variable duration CS, timing should not occur during the CS, only during the ITI (when the ITI is fixed). So, the comparison of fixed and random ITI allows for an examination of timing, independent of the level of conditioning that may occur to the CS.

2. Materials and methods

2.1. Animals

Thirty-two male Sprague Dawley rats (Taconic Laboratory, Germantown, NY), age 61–64 days at the beginning of the experiment were housed individually in a colony room on a 12:12 light–dark cycle (lights off at 08:45 h). Dim red lights provided illumination in the colony room and the testing room. The rats were fed a total daily ration of 15 g consisting of 45-mg Noyes pellets (Improved Formula A) that were delivered during the experimental session and additional PM1 FormuLab 5008 food that was given in the home cage shortly following the daily sessions. Water was available ad libitum in both the home cages and experimental chambers.

2.2. Apparatus

Twelve chambers (internal dimensions: $25 \times 30 \times 30$ cm), each located inside of a ventilated, noise-attenuating box ($74 \times 38 \times 60$ cm), comprised the experimental apparatus. Each chamber was equipped with a food cup, a water bottle, and a speaker. A magazine pellet dispenser (Med Associates, Model ENV-203) delivered 45-mg Noyes (Improved Formula A) pellets into the food cup. Each head entry into the food cup was transduced by a LED-photocell. The water bottle was mounted outside the chamber; water was available through a tube that protruded through a hole in the back wall of the chamber. The speaker for delivering white noise was above and to the left of the water tube. Two Gateway 486 DX2/66 computers running the Med-PC Medstate Notation Version 2.0 (Tatham and Zurn, 1989) controlled experimental events and recorded the time at which events occurred with 10-ms resolution.

2.3. Procedure

Four procedures, diagrammed in Fig. 1, were delivered to four different groups of rats ($n = 8$). Each procedure contains two events: a CS (a 70 dB white noise) and a US (one 45-mg Noyes food pellet). Each procedure is characterized by two time-event lines, one for the delivery of the CS (upper time-event line) and one for the delivery of the US (lower time-event line). For each procedure, the time of occurrence of the CS is indicated

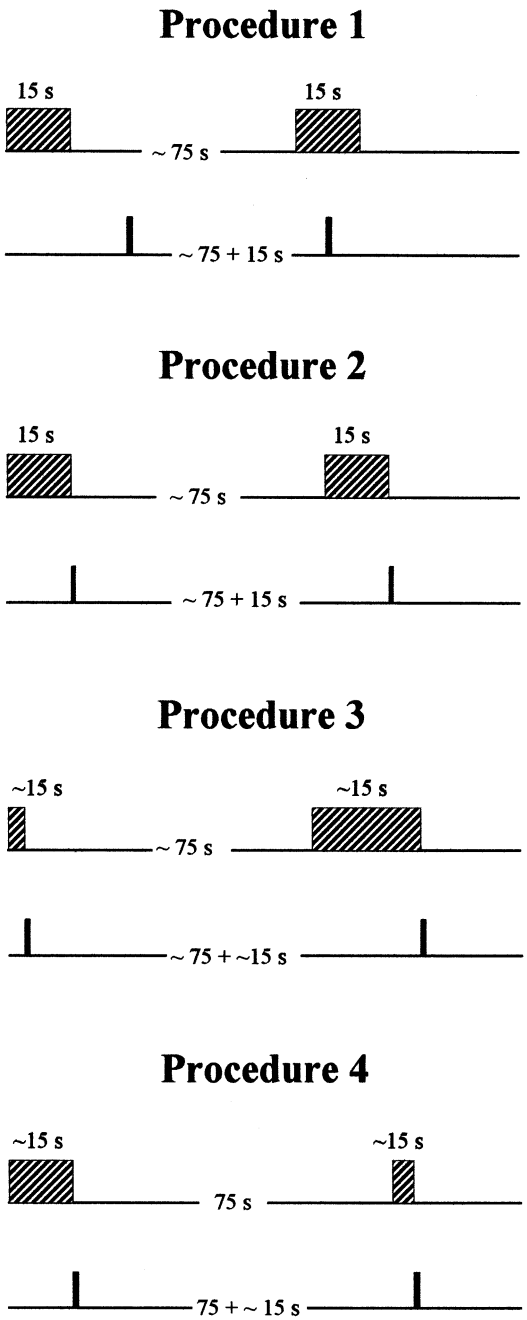

Fig. 1. A diagram of the four procedures that were delivered to separate groups of rats. Each procedural diagram contains two time-event lines, one for the delivery of signals (CSs) and one for the delivery of reinforcers (USs). The CS is denoted by a striped bar and its duration is indicated above the bar. The US is denoted by a dark bar. The duration between successive signals or reinforcers is indicated on the time-event line.

by a striped bar and the time of occurrence of the US is indicated by a dark bar. Because the CS duration in some procedures was variable, the striped bars are sometimes of different widths. The time between successive events (CS onset to CS termination, CS termination to next CS onset, and US to US) is indicated on each time-event line. For all of the time intervals in the figure, a tilde denotes that the interval was an exponentially distributed random waiting time; intervals without a tilde were of fixed duration. Sometimes, intervals were obtained by combining a random waiting time and a fixed time. For example, $\sim 75+15$ s (see Procedure 1) indicates that the interval was a random waiting time with a mean of 75 s added to a fixed time of 15 s. The random 75 s portion of the interval would have an expected mean of 75 s, an expected median of 52 s, and, given a large sample, an expected range (comprising 99% of the possible waiting times) of 0.38–397.37 s (Evans et al., 1993).

Procedure 1 was a control procedure in which the CS and US were delivered independently of one another. The CS was a fixed duration of 15 s, the interval between the termination of one CS and the onset of the next CS was a random waiting time with a mean of 75 s (~ 75 s), and the US to US interval was $\sim 75+15$ s. Procedure 2 was a standard delay conditioning procedure in which the US always occurred at the time of CS termination and the CS was 15 s in duration. Thus, Procedures 1 and 2 differed only in the arrangement of food—in Procedure 1 food was at random with respect to CS delivery, whereas in Procedure 2 food always occurred at the end of the CS. Procedure 3 was the same as Procedure 2 except that the duration of the CS was a random waiting time (~ 15 s) rather than a fixed duration (15 s). Procedure 4 was the same as Procedure 3 except that the duration between the termination of one CS and the onset of the next CS was fixed 75 s instead of a random waiting time (~ 75 s). So, the successive pairs of procedures (Procedure 1 vs. Procedure 2, Procedure 2 vs. Procedure 3, Procedure 3 vs. Procedure 4) differed from one another in only one aspect. Procedures 1 and 2 were delivered for 10 daily sessions lasting 2 h

each; Procedures 3 and 4 were delivered for six daily 2-h sessions. The conditioned response that was measured was the head-in-magazine response.

2.4. Data analysis

Both a 'response strength' and a 'response timing' measure were examined. The response strength measure was a discrimination ratio. The numerator of the ratio was the response rate in the CS (responses/min) which was obtained by summing the total responses that occurred during the CS and dividing by the total number of minutes of the CS in each session. The denominator of the ratio was the sum of the response rates in the CS and non-CS (ITI) periods. The response rate in the ITI was obtained by totaling the responses that occurred in the non-CS period and dividing by the total number of minutes of the non-CS period in each session.

The response timing measure was the local rate (responses/min) of responding as a function of time since CS onset. Responses were examined in each 1-s interval from one CS onset to the next CS onset. The time of occurrence of each response was calculated relative to the time of the preceding CS onset. The number of occurrences of each of these CS–CS intervals (opportunities) was also calculated. The ratio of number of responses in each interval to the number of opportunities in each interval, multiplied by 60, is the responses per minute as a function of time since the previous CS.

Statistical analyses were conducted using two-tailed t-tests, with a preset alpha level of 0.01.

3. Results and discussion of empirical data

3.1. Food at end versus food at random

The first comparison, displayed in Fig. 2, was between Procedures 1 and 2. The procedures differed only with regard to whether food was at the end of the CS (Procedure 2) or at random times with respect to the delivery of the CS (Procedure 1). In Procedure 1, the expected time to food was independent of the occurrence of the CS; in Procedure 2, the expected time of food was 15 s after CS onset.

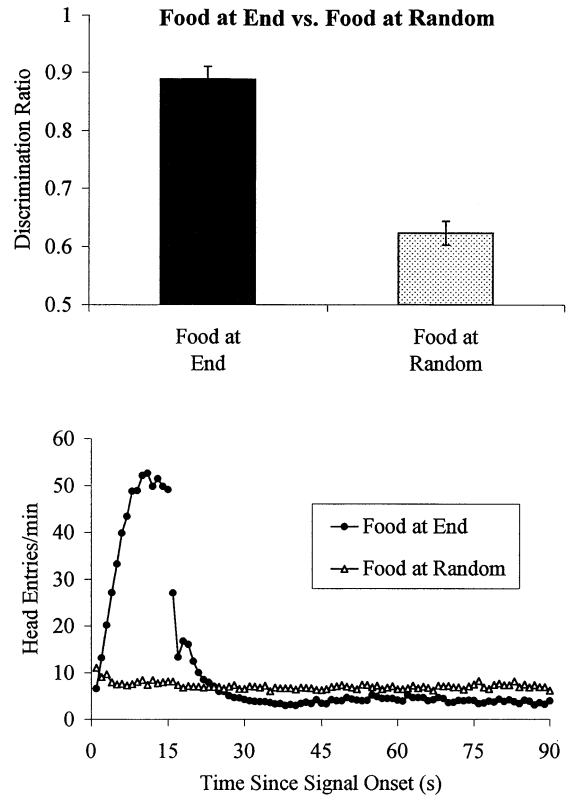

Fig. 2. A comparison of the effects of delivery of food at the end of the signal (Procedure 2) or with food at random times (Procedure 1). The top panel displays the discrimination ratios for the two conditions, averaged over the last half of training. The bottom panel displays the mean head entries per minute as a function of time since signal onset for the two conditions. The time of signal termination (15 s) is indicated by a break in the functions.

In Fig. 2, and subsequent results figures, the top panel displays the discrimination ratio and the bottom panel displays the local response rate functions; both dependent measures were obtained at steady state (the last half of training). The discrimination ratio revealed that food delivered at the end of the CS resulted in significantly greater responding to the CS (relative to responding in the non-CS period) than delivering food at random ($t(14) = 10.7$). This is the typical result when delay conditioning procedures (Procedure 2) are contrasted against control procedures (Procedure 1).

The local response rate measure revealed that when food was delivered at random with respect to the CS, the response rates were fairly constant across the CS and non-CS periods. When food was delivered at the end of the CS, there was an increasing response rate function that peaked shortly before the time of food delivery and decreased after food delivery. A linear function was fit to the normalized response rate functions (the proportion of the maximal rate) during the CS and revealed a significant difference between the slopes in the two conditions [$t(14) = 5.4$]. The slopes in the food at end condition averaged $+0.079$, whereas the slopes in the food at random condition averaged -0.009. Thus, the temporal gradients reveal that rats in the food at end procedure learned to: (1) respond more during the CS than during the non-CS period, which is captured by the discrimination ratios; and (2) respond more late in the CS near the time of upcoming food than early in the CS when the upcoming food delivery is temporally remote, which is captured by the shape of the local response rate functions.

3.2. Fixed versus random CS duration

Fig. 3 displays the comparison between a fixed duration CS (Procedure 2) and a random duration CS (Procedure 3) in appetitive conditioning. Because the mean CS–US and US–US intervals were the same, the presence of the CS (vs. the background) was equally predictive of food in the two groups. But, the local predictability of food differed between the procedures. In the fixed duration CS procedure, the CS not only predicted that food would occur, but also predicted that food would occur 15 s after the onset of the CS. In the random duration CS procedure, the CS only predicted that food would occur, but the time since CS onset provided no information regarding the actual time of food delivery.

As the discrimination ratios in Fig. 3 show, the fixed and random duration CSs were discriminated at a high level, indicating that both fixed and random CS durations were capable of supporting high levels of conditioning. The two procedures were not statistically different at the preset alpha level of 0.01 ($t(14) = 2.2$). Because the confidence intervals were very small in both groups, the small mean difference did yield a t-value that would have passed at an alpha level of 0.05.

An examination of the local response rates revealed a profound difference between the two conditions. The fixed duration CS resulted in an increase in the response rate during the CS. The random duration CS resulted in a sharp increase 1 s after CS onset followed by a decline in responding, and then a relatively constant rate of responding (the response rate remained relatively

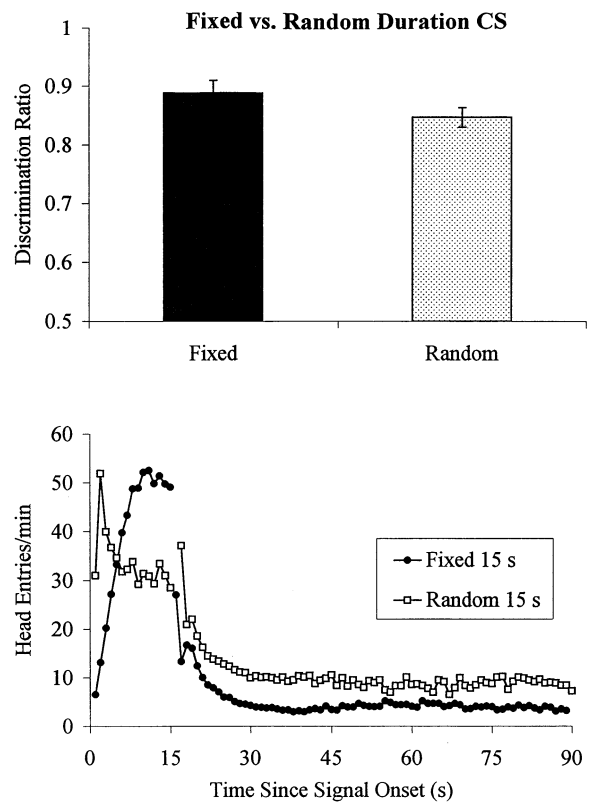

Fig. 3. A comparison of fixed (Procedure 2) and random CS durations (Procedure 3). The top panel displays the discrimination ratios for the two conditions, averaged over the last half of training. The bottom panel displays the mean head entries per minute as a function of time since signal onset for the two conditions. The time of signal termination in the fixed duration (15 s) and the average time of signal termination in the random duration (15 s) is indicated by a break in the functions.

constant during the CS beyond the mean of 15 s—data not shown). A linear fit of the normalized local response rate functions for each rat revealed that the slope of the best-fitting line was significantly more positive in the fixed CS condition than in the random CS condition ($t(14) = 5.0$). In summary, although the discrimination ratios revealed only a subtle difference between fixed and random duration CSs, the temporal gradients were clearly different in form. The differences in the form of the gradients are consistent with those reported by Libby and Church (1975).

3.3. Fixed versus random ITI

A comparison of the discrimination ratios in the fixed and random ITI conditions, revealed no statistical difference ($t(14) = -0.3$); robust levels of conditioning were evident in both conditions. So, much like the CS variability results, the variability of the ITI did not have an appreciable effect on discrimination accuracy.

The local response rate functions did, however, reveal a substantial difference between the conditions. The most striking difference was in the non-CS period. The random group produced a relatively constant response rate from 30 to 90 s after signal onset, with a mean of around ten head entries/min; the mean slope of the best-fitting straight line was 0.000 over this observation period. On the other hand, the fixed group produced an increasing response rate function over the same observation interval; the mean slope of the best-fitting straight line was $+0.010$. There was a significant difference in the slopes of the best linear fits between the conditions ($t(14) = 6.7$). The increasing response rate in the fixed group is indicative of temporal anticipation of the next reinforcer (see Church et al., 1991 for a similar result in a peak procedure).

3.4. Summary of results

The paired comparisons of the four procedures revealed that the discrimination ratios reflected variations in the global, relative rate of food delivery in the CS and non-CS periods. When there was a difference in global predictability (food at end vs. food at random), then the discrimination ratios were different, and when there were no differences in global predictability (fixed vs. random duration CS, fixed vs. random duration ITI), then the discrimination ratios were similarly high.

The local response rate functions, on the other hand, reflected both local and global predictability. The shape of the local response rate functions revealed differences in the local predictability of food. For example, response rates increased during the CS (or ITI) when the CS (or ITI) was fixed in duration. The local rate functions also disclosed the effect of global predictability: the relative height of the function in the CS and non-CS periods provided the same index as the discrimination ratios. So, if any theory of timing or conditioning could adequately recreate the local rate functions, then the discrimination ratios would automatically be reproducible.

The overall pattern of results, with global predictability of food affecting the global distribution of responses and local predictability of food affecting the local distribution of responses, is consistent with Gibbon and Balsam's (1981) account of conditioning (see also Jenkins et al., 1981). Specifically, they argued that two timing mechanisms are involved in conditioning: a gross timing mechanism and a fine timing mechanism. The gross timing mechanism is sensitive to the global rate of reinforcement in the CS and background. When the ratio of the expectancy of reinforcement in the background and CS periods (the c/t ratio) exceeds some threshold, then responses will occur during the CS. The gross timing mechanism therefore controls the decision of whether responses will occur in the presence of a stimulus. Once responding is well established by the gross timing mechanism, then the fine timing system is initiated. The fine timing system is simply Scalar Timing Theory, already described above, and it controls the decision of when responses will occur in the presence of a stimulus.

The Gibbon and Balsam account is a dual process model in which conditioning (the gross timing mechanism) and timing (the fine timing mechanism) are separate and are initiated serially. The notion that the timing system is much slower

than the conditioning system is a common view, as represented by a comment by Sutton and Barto (1990) "… when animals are presented with a very long-duration CS, followed by the US, they eventually learn to respond differentially to the earlier and later portions of the CS." Interestingly, there is very little evidence in support of this notion. Gibbon et al. (1980) found that the acquisition of responding (conditioning) emerged much earlier than the acquisition of a well established temporal gradient during the CS (timing). However, they did not attempt to directly compare the emergence of these two features of responding by using a common metric. On the other hand, Holder and Roberts (1985) conducted a systematic investigation of the emergence of timing and conditioning and concluded that the two features of responding emerged almost simultaneously. While further investigation will undoubtedly resolve the differences in empirical outcomes of these two studies, the Holder and Roberts investigation raises the possibility that these two influences on responding may arise together, perhaps even through a single process. As shown in the following section, a hybrid model that requires only one process for the emergence of response strength and timing can provide a reasonably good qualitative fit to the results in Figs. 2–4.

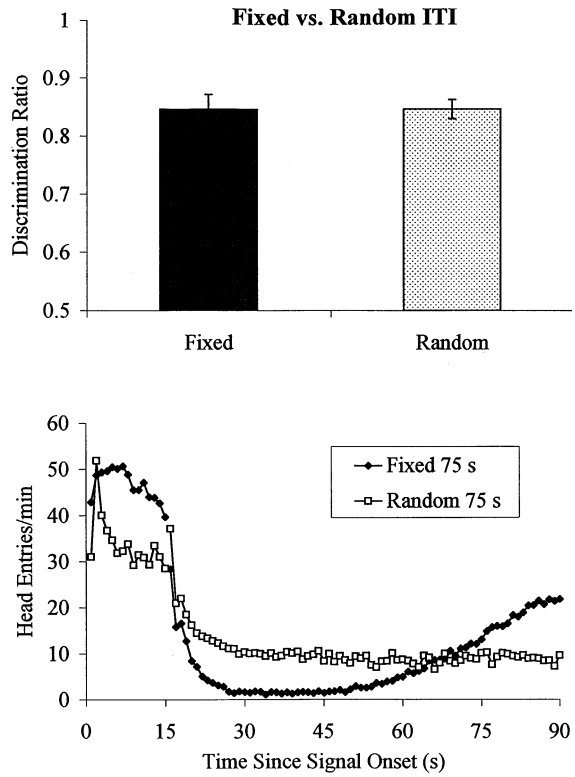

Fig. 4. A comparison of fixed (Procedure 4) and random ITI durations (Procedure 3). The top panel displays the discrimination ratios for the two conditions, averaged over the last half of training. The bottom panel displays the mean head entries per minute as a function of time since signal onset for the two conditions. The average time of signal termination (15 s) is indicated by a break in the functions.

4. Model simulations

4.1. Specification of the model

The Behavioral Theory of Timing (Killeen and Fetterman, 1988) was developed to account for the time of occurrence of responses that did not require the assumption that the animal was an information-processing system with modules for temporal perception, memory, and decision. Instead, the basic assumption was that the animal changed from one behavior state to another and that, when reinforcement occurred, the behavior state at that instant was strengthened. Machado (1997) provided a precise statement of a version of this theory that accounted for behavior in many timing experiments (hereafter referred to as Machado's Behavioral Theory of Timing, MBT). Despite the behavioral origins of the theory, it may be described in the same information-processing framework that has been used to describe Scalar Timing Theory.

There are three parts of the theory: behavioral states, learned associations, and a response rule. The behavioral states may be considered to be the perceptual representation of a time interval, $X_{\cdot,t}$, that is initiated by an event (a change in the status of a stimulus); the learned associations serve as a memory of the time of reinforcement relative to some event, $V_{\cdot,t}$; and the response rule is the basis for the decision whether or not to respond at any given moment in time, R_t.

Any event may serve as a time marker: stimulus onset, stimulus termination, reinforcement, or re-

sponse. When a time marker occurs, a cascade of functions is initialized; the first function is set to 1 and the others to 0. Then, as time progresses, a family of n gamma functions emerge as shown in the top panel of Fig. 5; n is a parameter of the model. This figure is based on Eq. (1), in which t is the time since the time marker, i is the individual function number, and λ controls the rate of change. At any given time (T) since the time marker, the perceptual representation of that time is $X_{:,T}$, where this is a vector of n activation heights, one for each function ($1 \leq i \leq n$). Thus, time is not represented by an amount (as in Scalar Timing Theory) but as a pattern of heights of the individual activation functions.

$$X_{i,t} = X_{i,t-1} - \lambda X_{i,t-1} + \lambda X_{i-1,t-1} \quad \text{for } i > 1 \tag{1}$$

The strength weights for the individual functions (which represents the memory for reinforcement), $V_{i,t}$, increase at the time of reinforcement and decrease at all times. The updating equations for the strength weights of the individual functions ($1 \leq i \leq n$) are:

$$V_{i,t} = V_{i,t-1} + \alpha[X_{i,t-1}*(1 - V_{i,t-1})] \tag{2a}$$

$$V_{i,t} = V_{i,t-1} + \beta[X_{i,t-1}*(0 - V_{i,t-1})] \tag{2b}$$

Eq. (2a) only applies if reinforcement occurs on time step t; Eq. (2b) applies at all time steps. Note that $X_{i,t}$ is the height of the ith activation function, and $V_{i,t}$ is the corresponding strength weight of the ith function. The α/β ratio is an index of the relative impact of reinforcement.

An example of the strength weights obtained from Eqs. (2a) and (2b) is shown in the middle panel of Fig. 5. This particular distribution of strength weights was obtained from a procedure in which reinforcement was delivered 15 s after a time marker; the distribution is made up of the asymptotic strength weights of each function on the last time step of a 1000-m simulation session. At asymptote, the strength of the memory for reinforcement, $V_{i,t}$, is relatively constant for each function. As seen in the middle panel of Fig. 5, Functions 8 and 9 received the most weight. Referring back to the pattern of activation in the top panel, one can see that Functions 8 and 9 were highly active at 15 s after the marker (indicated by the arrow). Thus, the amount of activation of each function at a particular time reflects the eligibility of that function to be modified by reinforcement. The consequent distribution of strength weights reflects which functions were most active at the time of reinforcement.

The decision to respond at time t is based on the sum of the products of the height of each

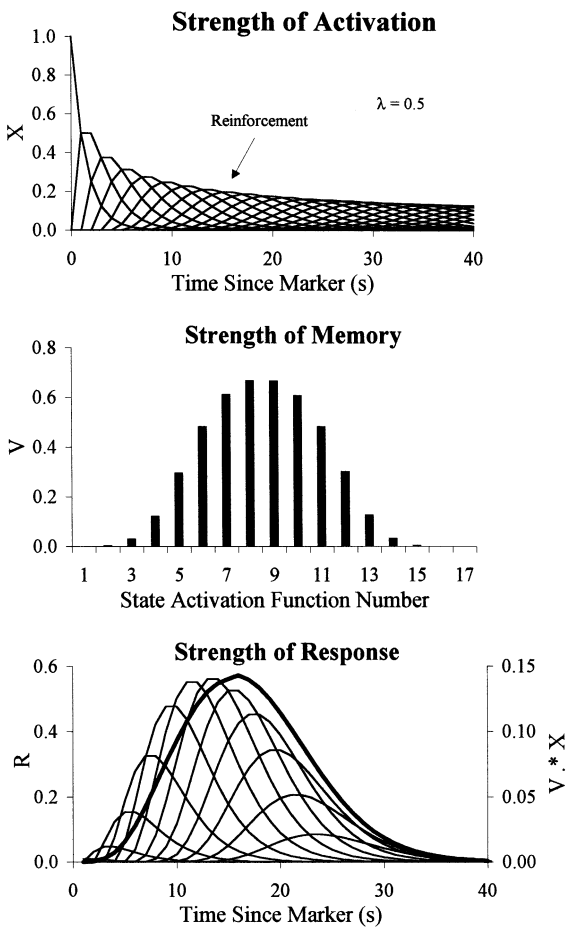

Fig. 5. Top panel: Strength of activation (X) of the cascade of gamma functions as a function of time since a marker with a lambda value of 0.5. A possible time of reinforcement is indicated by the arrow. Middle panel: Strength weights of the memory for reinforcement (V) associated with each of the first 17 state activation functions. Bottom panel: the family of functions generated by the product of the individual state activation functions (X) and their associated strength weights (V); the thick line is the sum of the individual products, R.

activation function ($X_{i,t}$) and its related strength weight ($V_{i,t}$) at that time, scaled by a constant η. Machado (1997) used this deterministic quantity, R_t, to map onto response rate measures, but indicated that responses could be generated stochastically from the strength of R_t.

$$R_t = \eta * \Sigma X_{i,t} V_{i,t} \quad (3)$$

The individual products of the activation functions and their associated weights are displayed in the bottom panel of Fig. 5 (thin lines), along with the sum of the products (heavy line). Functions 8 and 9, which received the most weight and were most active at the time of reinforcement, are the two highest individual functions. The sum of the weighted activation functions forms a gradient that has a maximum near the time of reinforcement.

4.2. Implementation of the model

The MBT model described above was implemented for purposes of simulation of the results displayed in Figs. 2–4. Machado's (1997) specification of the model included only a single time marker, but two time markers were required for simulation of the above results: one for stimulus onset and one for reinforcement. In our implementation of the MBT model, these two time markers initiated separate activation and strength processes. Eqs (1–3) were used with the time of occurrence of the most recent CS as the time marker to produce X_{CS}, V_{CS}, and R_{CS}. These equations were also used with the reinforcement as the time marker to produce X_{US}, V_{US}, and R_{US}.

The strength of the response at time t, R_t, was converted into all-or-none responses. A response was assumed to occur at time t whenever R_{CS} or R_{US} was greater than a uniform random number between 0 and 1.

The simulation program was written in a modular form in Matlab, version 5.0 (The MathWorks, Inc., Natick, MA). The calling program was:

```
1. iprocedure_procedurename
     % Initialize the procedure
2. imodel_MBT
     % Initialize the model
3. for session_clock = 0:session_duration
       4. procedure_procedurename
            % Run the procedure for 1 s
       5. model_MBT
            % Run the model for 1 s
6. end
7. record
     % Record the data
```

The procedure determines the times of occurrence of stimuli and reinforcements, and it receives information about the times of occurrence of responses from the model. The model determines the times of occurrence of responses, and it receives information about the occurrence of stimuli and reinforcements from the model (see Church, 1997 for further description of this modular approach.) By using a modular structure, as demonstrated by the Matlab code above, one can easily change the procedure (Lines 1 and 4), without having to modify the original code associated with the model (Lines 2 and 5). This separation of model and procedure provides confidence that the same model of the animal's processing system is being used for different procedures.

Because the basic data structure of Matlab is the matrix, the code is easy to read and write. For example, the Matlab code for Eq. (2a), which updates the strength of the CS (V_{CS}) at the time of reinforcement, is: $V_{CS} = V_{CS} + \alpha * X_{CS} .* V_{CS}$. The '.*' indicates that there was array multiplication (multiplication of corresponding elements) of the vectors X_{CS} and V_{CS}, which contained the individual values of activation and strength for the n activation functions.

4.3. Selection of parameters

The session duration of 1000 min resulted in the delivery of approximately 600 stimuli and reinforcers in each run of the simulation, an amount comparable to that received by the rats. The initial parameters were chosen on the basis of parameters that were used by Machado (1997) and were then modified slightly to improve the quality of fit to the experimental data. Only three runs of the simulation (with variations in parameters) were required in order to obtain parameters that provided good qualitative fits of the experi-

mental data. The simulation results, which are displayed in the following section, were obtained on the third run.

Because the mean signal to reinforcer and reinforcer to reinforcer intervals were the same in all four procedures, the same parameter values could be used for all of the simulations. The parameter values were as follows: $n = 60$, $\lambda = 0.5$, $\alpha = 0.01$, $\beta = 0.2$, and $\eta = 1$. The initial values of V_{CS}, V_{US}, R_{CS}, and R_{US} were set to 0. The activation functions were initialized so that the activation of the first function at the first time, relative to both the CS and the US marker, was set to 1 and all other functions were 0.

4.4. Simulation results

4.4.1. Food at end versus food at random

The results of the MBT simulation on the Procedures 1 and 2 are displayed in Fig. 6. The organization and content of Fig. 5 is the same as the actual data that were displayed in Fig. 2. There are two striking features of the simulation output. First, the discrimination ratios are ordered correctly, but are of much lower magnitude than the actual data, particularly for the food at end procedure. However, the temporal gradients were quite similar to the actual data. The food at random procedure resulted in a constant rate of responding that was approximately 10–15 responses/min (the rats responded at about ten responses/min). The food at end procedure resulted in an increasing rate of responding during the signal and then a decreasing rate of responding after signal termination, falling to an ultimate baseline rate of around ten responses/min. The major difference in the simulation data (which contributed to both the discrimination ratios and the shape of the response rate functions) was that the model produced more responding after signal termination than the rats—responding was both of greater magnitude and persisted for a longer duration. The higher responding after signal termination by the model contributed to a lower overall discrimination ratio.

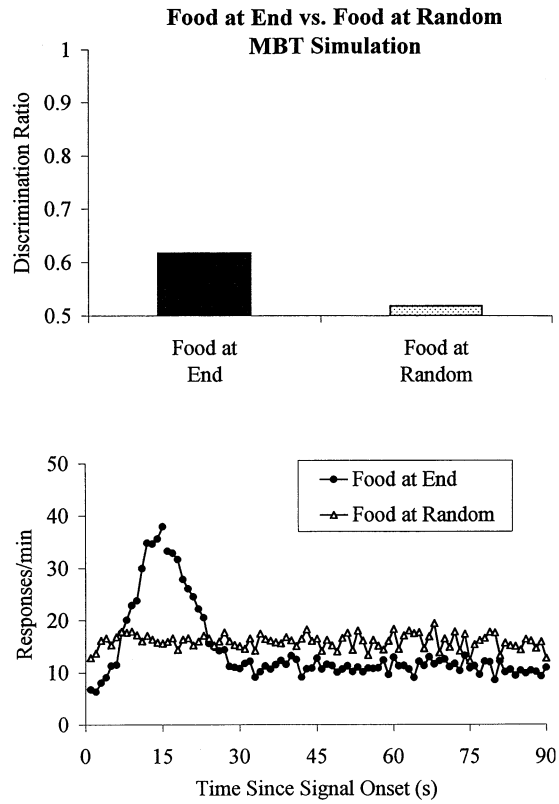

Fig. 6. Simulation results for the comparison of food at the end of the CS (Procedure 2) versus food at random times (Procedure 1).

4.4.2. Fixed versus random CS duration

The results of the simulations on Procedures 2 and 3 are displayed in Fig. 7 (these results can be compared to the empirical data portrayed in Fig. 3). The results of the MBT simulations were qualitatively impressive in that they captured the primary difference between fixed and random duration CSs: the response rate during the CS was rising when the CS was a fixed duration and was relatively constant (after an initial sharp increase and decrease) when the CS was a random duration. The discrimination ratios produced by MBT simulations, for the two procedures were similar in magnitude. This was also observed in the data, although the discrimination ratios produced by the model were substantially lower than those achieved by the rats. The low discrimination ratios appeared to be due to greater responding in the nonsignal period by the model than the rats.

4.4.3. Fixed versus random ITI

The MBT simulation results on the fixed versus random ITI procedures are displayed in Fig. 8. Again, the qualitative features of the temporal gradients of the model's output were similar to the qualitative features of the rat's output (see Fig. 4). In particular, the responding during the CS increased sharply during the first 2 s of the CS duration and then was steady or modestly decreasing. Responding during the ITI decreased gradually in the Random group, but increased in the Fixed group in the later portion of the ITI. Thus, the model demonstrated the key feature that was also present in the rat's data: temporal anticipation of the upcoming reinforcer during the fixed ITI.

The model's performance, as measured by the discrimination ratios, was qualitatively different from the performance of the rats. The discrimination ratios produced by the model were lower than the discrimination ratios produced by rats. And, the discrimination ratios in the fixed ITI were higher than in the random ITI, a trend that was not present in the rat data. There was again evidence of too much responding in the ITI, particularly in the random ITI procedure, which resulted in a low discrimination ratio.

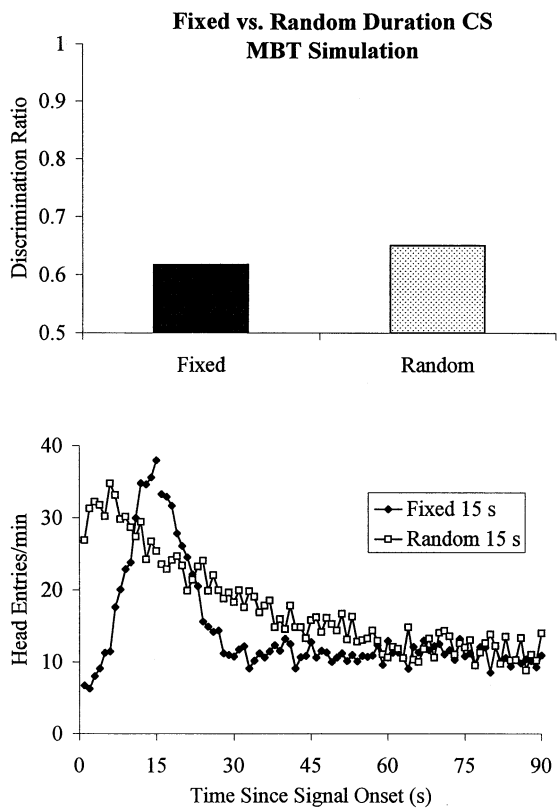

Fig. 7. Simulation results for the comparison of a fixed (Procedure 2) and random (Procedure 3) duration CS.

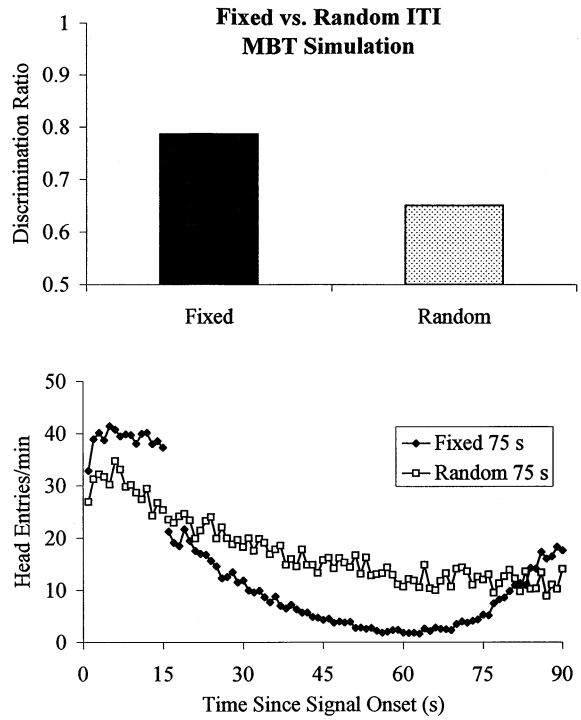

Fig. 8. Simulation results for the comparison of a fixed (Procedure 4) and random (Procedure 3) duration ITI.

5. General discussion

The MBT model (Machado, 1997) accounted qualitatively for performance of rats on four procedures that can be viewed either as conditioning or timing procedures. The global response rates were shown as discrimination ratios, and the local rates were shown as gradients of responding as a function of time since stimulus onset. The MBT model provided a basis for understanding how the global rate of reinforcement controls the global rate of responding, and how the local expected

time to reinforcement controls the local rate of responding. Quantitatively, the model's performance (shown in Figs. 6–8) was somewhat different from the performance of the rats (shown in Figs. 2–4). The major problem was that the model yielded considerable generalization of responding in the ITI, which resulted in low discrimination ratios, even in the procedures in which fairly optimal conditioning parameters were employed (e.g. Procedure 2). This failure of the model appears to be related to the asymmetrical generalization gradients that result from the cascade of gamma functions, an inherent characteristic of the function form of the perceptual representation of time.

We have not as yet conducted a broad search of the parameter space with this model, so improvements in the quantitative fit of the data would undoubtedly be possible. The parameters that could be modified are: λ, the intensity parameter of the gamma distribution that determines the rate of change of the activation functions, the α/β ratio, which determines the relative impact of reinforcement, and η, the constant that modifies the strength of the response tendency. All of these parameters have effects on the shape of the response rate functions. For example, modifications of the α/β ratio affects the height and width of the response gradient, but does not alter the asymmetrical shape of the gradient.

Another problem with the perceptual representation of time in the MBT model is that the cascade of gamma functions provides a basis for timing only within a restricted range—beyond some duration the number of functions with available activation decreases toward zero. This is due to the fact that activation continues to flow out of the system, after the maximum number of functions has been reached. For the present simulations, with a value of λ that provided adequate response gradients during the CS, 60 functions were required in order to time both the short duration CS and the longer duration US–US interval.

As noted in the Introduction, real-time conditioning models also provide a basis for predicting the effects of classical conditioning procedures on response strength and timing measures of performance. Real-time conditioning models are usually considered to be variants of state-based conditioning models, with time explicitly represented. The equations of real-time conditioning models, however, may also be described in terms of modules for perception, memory, and decision. In Fig. 5, the three parts of Machado's version of the Behavioral Theory of Timing were labeled 'strength of activation,' 'Strength of memory,' and 'Strength of response'. Real-time conditioning models also contain the same three parts.

An example of a real-time conditioning model is the complete-serial-compound version of the time-derivative model (TDM) of Pavlovian conditioning (Sutton and Barto, 1990). The perceptual representation of a stimulus, such as a noise that lasts for 15 s, is equivalent to a series of different component stimuli each lasting for a short interval of time (t) that are activated in a serial fashion (see Fig. 9, top panel). Eq. (4) specifies the perception ($X_{i,t}$) of each of these components (i) at time t; $\kappa_{i,t-1}$ denotes the on-off status of stimulus component i (1 or 0) on the previous time step ($t-1$). At the onset of the stimulus (T), all elements are initialized at 0 ($X_{.,T}=0$). Stimulus components are successively activated, one at each time step. Once an element has been activated, its strength decays exponentially, with the rate of decay determined by λ.

$$X_{i,t} = X_{i,t-1} - \lambda X_{i,t-1} + \lambda \kappa_{i,t-1} \tag{4}$$

The memory representation in the TDM model is a vector of strength weights, with a different strength weight calculated for each function at each time step. Eq. (5) specifies the strength ($V_{i,t}$) of each stimulus component (i) at time t. The variable, U_t codes the presence ($U_t=1$) or absence ($U_t=0$) of reinforcement at time t; α, β, and γ are weighting parameters.

$$V_{i,t} = V_{i,t-1} + \beta(U_t + \gamma R_t - \mathrm{R}_{t-1})\alpha X_{i,t} \tag{5}$$

Changes in strength occur at CS onset, CS termination, during the CS, and at the time of US presentation. At CS onset, $\Delta V_{i,t} = \beta\gamma R_t \cdot \alpha X_{i,t}$. During the CS, $\Delta V_{i,t} = \beta(\gamma R_t - R_{t-1})\alpha X_{i,t}$. At CS termination, $\Delta V_{i,t} = -\beta R_{t-1} \cdot \alpha X_{i,t}$. The US leads to $\Delta V_{i,t} = \beta(1 + \gamma R_t - R_{t-1})\alpha X_{i,t}$. Thus, CS onset and US presentation lead to increments in V,

whereas CS termination leads to a decrement in V. Notice that every change in strength includes the multiplicative factor, $\alpha X_{i,t}$, which means that the level of activation of the various components will always affect the strength weights. Thus, the parameter λ, which determines the shape of the activation functions produces important effects on the strength weights. The other factor which is involved in instances where increments in V occur is γ, which is a discount parameter that reflects imminence weighting of the upcoming US. The primary effect of γ is to determine the shape and peak rate of the response rate function, R, with smaller values of γ producing more accelerated response rate functions, but with a lower peak rate.

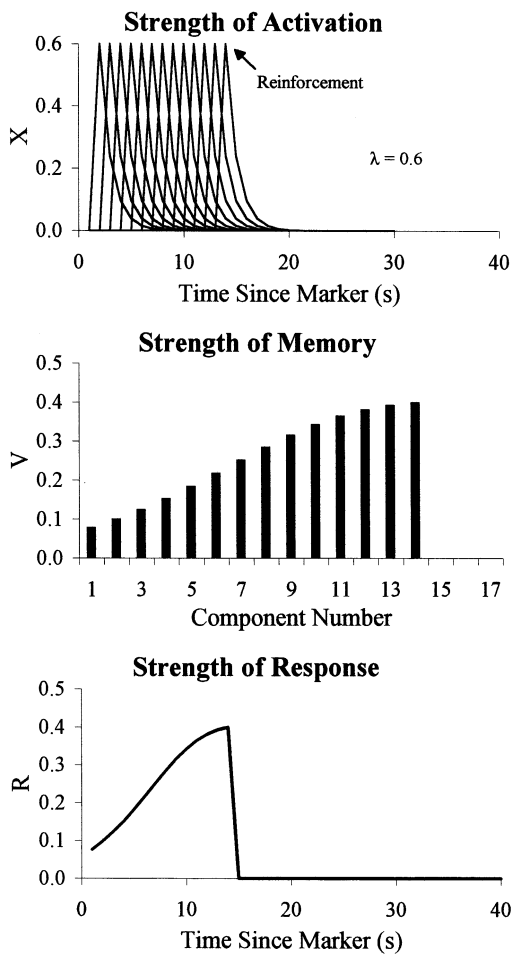

Fig. 9. Top panel: Strength of activation (X) of the serially activated stimulus components as a function of time since a marker with a lambda value of 0.6. A possible time of reinforcement is indicated by the arrow. Middle panel: Strength weights of the memory for reinforcement (V) associated with each of the first 14 stimulus components. Bottom panel: the strength of response, R, as a function of time since a marker.

The middle panel of Fig. 9 displays a set of strength weights obtained after training for 300 min (178 trials) with Procedure 1 (see Fig. 1), using 1-s time steps, $\lambda = 0.6$, $\gamma = 0.99$, $\beta = 1.0$, and $\alpha = 0.05$. The strength weights were obtained at the end of training, during the last time step of the last CS in the simulation. In general, the components that are activated nearer to the time of reinforcement receive greater strength weights, because the shape of the memory structure is influenced by the pattern of activation at the time of reinforcement. However, unlike MBT, the shape of the strength gradient does not directly represent the strength of the activation functions. This is due to the discounting parameter, γ, which reduces the weight of the current information about the imminence of the US. Also notice that only those functions that were active prior to the time of reinforcement receive any strength; there is no generalization after reinforcement. This is due to the fact that components are no longer activated once the CS is terminated. The TDM model can generate strength weights after CS termination if the implementation of the model includes a cascade that is initiated when the CS is terminated (Moore and Choi, 1997).

The response is produced by the variable R_t, which is the sum of the strength weights of all active components (see Eq. (6)). Because only one component is active at any time step, R_t is the strength of the active component. The bottom panel of Fig. 9 displays the strength of response obtained in the simulation with Procedure 1. Because the strength of response is determined directly by the strength of the active component, the response gradient is of the same shape as the strength weights, with an increasing response gradient that peaks 1 time step before the US.

$$R_t = \Sigma V_{i,t}\kappa_{i,t} \text{ if } R_t < 0, \text{ it is set to } 0 \quad (6)$$

A comparison of MBT and TDM involves contrasting their perceptual representations, memory

representations, and response strengths (see Figs. 5 and 9).

The perceptual representation of time in both models consists of a cascade of functions (see Eqs. (1) and (4)) that can be activated by any event (e.g. CS onset, CS termination, US occurrence). The functions used in these two theories differ in shape, but they serve to cover a range of physical times with multiple values. The decay process in the difference equations is the same, but the incremental process differs. Increments of activation in MBT come from the previous function; increments in TDM come from the on–off status of the element on the previous time step. This difference results in a gradual increase in activation in the MBT model vs. a sharp increase in activation in the TDM model. Also, in MBT, activation continues to flow through the system after the CS is terminated, whereas, in TDM, new components are fired only as long as the CS is on.

The perceptual representations in these two models suffer from two major problems: they require a large number of functions for timing to occur over a wider scale (e.g. in the above implementation of MBT, 60 functions were required in order to concurrently time the 15 s CS and 90 s US–US intervals); and they do not easily produce the scalar property. One nice feature of MBT is that the activation heights become more similar at longer delays than at shorter delays, meaning that shorter times are more discriminable than longer times.

The memory representation in both models consists of a vector of strength weights. In both models, the memory representation is strongly related to the pattern of activation at the time of reinforcement, but the updating equations differ fundamentally (see Eqs. (2a), (2b) and (5)). There are three major differences in the learning rules. The first difference is in the conditions under which decrements in V occur. In MBT, strength is assumed to decrease at all times, except when the reinforcer occurs. So, strength is lost during the CS and during the ITI. In TDM, decreases in strength can only occur at CS termination; during the ITI $V_{i,t} = V_{i,t-1}$. Accordingly, the MBT model will produce extinction with more ease than the TDM model. For extinction to occur in the TDM model, the decrimental effect of CS termination must outweigh the incremental effect of CS onset, which means that γ must be small.

The second difference between the TDM and MBT learning rules is in the conditions under which increments in V occur. In both models, reinforcement increments strength. However, in TDM, strength is also incremented at the time of CS onset. The difference is due to the fact that the TDM model encodes the on-off status of each component of a given CS through the variable κ_i. The variable κ_i, serves to encode both the state and the passage of time since an event. On the other hand, the MBT model encodes only events. By encoding the state, the TDM model can produce stimulus competition effects (e.g. blocking, overshadowing). By encoding the event, the TDM model can produce temporal effects (e.g. response timing, CS–US interval effects, etc.). The MBT model can only account for temporal effects in its present form, but it could be extended to encode the state as well as the event by the addition of a variable such as κ_i.

Finally, a third difference in the learning rules is in the factors that produce the structure of memory. In both models, changes in strength are due to a combination of activation ($X_{i,t}$) and strength ($V_{i,t}$). However, in TDM, the discount factor (γ) produces important additional effects on the structure of memory

While the learning rules in the real-time models appear to be promising in their ability to accommodate many stimulus and temporal variables, they should be explored more thoroughly under a wider range of procedures from both the conditioning and timing traditions. For example, it is not clear that the real-time models are capable of generating the scalar timing property. Moreover, there has not yet been a serious attempt to produce good quantitative fits to the data. The major basis for evaluating a hybrid theory is the extent to which the predictions of the model correspond to the quantitative features of the behavior of animals. A quantitative fit is much stronger support than a qualitative one, especially if the parameters of the model are individually affected by different independent variables. Thus, a more rigorous application of the real-time theories may disclose the need for an alternative approach.

The output of the two models (strength of response) is similar (Eqs. (3) and (6)). In both cases, the response strength is determined by a multiplicative combination of the memory representation and current perception. So, in both models, the shape of the response rate function directly reflects the structure of the memory representation. However, in MBT, the current perception is the level of activation of each of the functions; in TDM, the current perception is simply the on–off status of each of the components.

Placing the finer differences in detail aside, the structure of the TDM real-time conditioning model is the same as the structure of the MBT timing model. There is no fundamental distinction that requires one of these models to be considered a timing model and the other to be considered a conditioning model. They may both be considered hybrid models that are designed to account for both global and local features of response rate data.

Many other hybrid models have been proposed, and several of them have been completely specified. Two examples are the Spectral Timing Model (Grossberg and Schmajuk, 1989) and the Multiple-oscillator Model (Church and Broadbent, 1990). The Spectral Timing Model contains three parts: Perception, memory, and decision. A stimulus produces a series of internal events leading to a cascade of gated signal functions that are different in form, but similar in function, to those shown in the top panels of Figs. 5 and 9. Reinforcement leads to the formation of a long-term memory trace. A response is a product of the current perceptual representation and the memory trace.

The Multiple-oscillator Model also contains three parts: Perception, memory, and decision. A stimulus initiates a bank of periodic functions, and the vector of half-phases of each of these functions at any given physical time serves as the perceptual representation of time. The memory of reinforced times is stored in a autoassociation matrix, and decision is based on the similarity of the current perception of time with a retrieved value from the matrix.

Other examples include Sutton and Barto (1981), Desmond and Moore (1988), Klopf (1988) and Gluck et al. (1990). The interested reader should consult these references for a description of the calculations involved in these theories. The major point here is that many theories, which may be referred to as conditioning or timing models all contain the same three-part structure of perception, memory, and decision.

None of these theories are without their faults. With a modular approach to the development of theory (Church, 1997), it may be possible to modify a single module of one of these theories that would result in a version of the theory that accounts quantitatively for many different measures of behavior under many different procedures, is neurally plausible in terms of its perceptual representation and memory structure, and requires a relatively small set of parameters.

For example, the perception of time in the TDM model may be replaced by a set of functions that increase more gradually before decaying. In general, the perception of time may be represented by a single function, a few functions, or by a large number of functions. A single increasing or decreasing function is sufficient to provide a unique perception for every physical time. An increasing linear function is used in Scalar Timing Theory (Gibbon et al., 1984). An alternative is a negatively accelerated increasing function beginning at stimulus followed by a negatively accelerated decreasing function beginning at stimulus termination (Sutton and Barto, 1981). Periodic functions such as those in the Multiple-oscillator model (Church and Broadbent, 1990) provide a unique perception for every physical time within a single period, and phase discrimination within a circadian rhythm has been used to account for the ability of animals to go to different places at different times of day (Krebs and Biebach, 1989; Sakside and Wilkie, 1994).

Of course, it is possible that no hybrid model of timing and conditioning can account for the myriad of results produced by the domain of procedures involving multiple stimuli and reinforcements. As noted earlier, Gibbon and Balsam (1981) have proposed that the principles responsible for the growth of associative strength are different from the principles responsible for the timing of responses, and that these two pro-

cesses operate serially. It has not been established whether two separate mechanisms are required, or whether a single learning mechanism represented by a hybrid model is sufficient.

In his article, 'Are theories of learning necessary', Skinner (1950) argued for the need for accurate information regarding the empirical determinants of behavior. Now there is extensive information available about the variables controlling behavior, and there is general consensus that theories of conditioning and timing are useful. But it is not clear that separate theories are necessary. A theory that predicted the time of occurrence of responses during an experimental procedure in which stimuli and reinforcements were presented would not need to be supplemented by another theory that predicted average performance.

Acknowledgements

This research was supported by a National Institute of Mental Health grant MH44234 and a National Research Service Award MH1169 to Brown University.

References

Barnet, R.C., Cole, R.P., Miller, R.R., 1997. Temporal integration in second-order conditioning and sensory preconditioning. Anim. Learn. Behav. 25, 221–233.

Barnet, R.C., Miller, R.R., 1996. Temporal encoding as a determinant of inhibitory control. Learn. Motiv. 27, 73–91.

Blazis, D.E.J., Desmond, J.E., Moore, J.W., Berthier, N.E., 1986. Simulation of the classically conditioned nictitating response by a neuron-like adaptive element: A real-time invariant of the Sutton-Barto model. Proceedings of the Eighth Annual Conference of the Cognitive Science Society, Lawrence Erlbaum Associates, Hillsdale, NJ.

Church, R.M., 1997. Quantitative models of animal learning and cognition. J. Exp. Psychol.: Anim. Behav. Process. 23, 379–389.

Church, R.M., Broadbent, H.A., 1990. Alternative representations of time, number, and rate. Cognition 37, 55–81.

Church, R.M., Meck, W.H., Gibbon, J., 1994. Application of scalar timing theory to individual trials. J. Exp. Psychol.: Anim. Behav. Process. 20, 135–155.

Church, R.M., Miller, K.D., Meck, W.H., Gibbon, J., 1991. Symmetrical and asymmetrical sources of variance in temporal generalization. Anim. Learn. Behav. 19, 207–214.

Cole, R.P., Barnet, R.C., Miller, R.R., 1995. Temporal encoding in trace conditioning. Anim. Learn. Behav. 23, 144–153.

Desmond, J.E., Moore, J.W., 1988. Adaptive timing in neural networks: The conditioned response. Biol. Cybernet. 58, 405–415.

Dickinson, A., 1980. Contemporary Animal Learning Theory. Cambridge University Press, Cambridge.

Dickinson, A., Mackintosh, N.J., 1978. Classical conditioning in animals. Ann. Rev. Psychol. 29, 587–612.

Durlach, P.J., 1989. Learning and performance in Pavlovian conditioning: Are failures of contiguity failures of learning or performance? In: Klein, S.B., Mowrer, R.R. (Eds.), Contemporary Learning Theories: Pavlovian Conditioning and the Status of Traditional Learning Theory. Lawrence Erlbaum Associates, Hillsdale, NJ, pp. 19–59.

Evans, M., Hastings, N., Peacock, B., 1993. Statistical Distributions, 2nd ed. John Wiley, NY, p. 170.

Gibbon, J., Baldock, M.D., Locurto, C.M., Gold, L., Terrace, H.S., 1977. Trial and intertrial durations in autoshaping. J. Exp. Psychol.: Anim. Behav. Process. 3, 264–284.

Gibbon, J., Farrell, L., Locurto, C.M., Duncan, H.J., Terrace, H.S., 1980. Partial reinforcement in autoshaping with pigeons. Anim. Learn. Behav. 8, 45–59.

Gibbon, J., Balsam, P., 1981. Spreading association in time. In: Locurto, C.M., Terrace, H.S., Gibbon, J. (Eds.), Autoshaping and Conditioning Theory. Academic Press, New York, pp. 219–253.

Gibbon, J., Church, R.M., 1990. Representation of time. Cognition 37, 23–54.

Gibbon, J., Church, R.M., Meck, W.H., 1984. Scalar timing in memory. Ann. N.Y. Acad. Sci. 423, 52–77.

Gluck, M.A., Reifsneider, E.S., Thompson, R.F., 1990. Adaptive signal processing and the cerebellum: Models of classical conditioning and VOR adaptation. In: Gluck, M.A., Rumelhart, D.E. (Eds.), Neuroscience and Connectionist Theory: Developments in Connectionist Theory. Lawrence Erlbaum Associates, Hillsdale, NJ, pp. 131–185.

Grossberg, S., Schmajuk, N.A., 1989. Neural dynamics of adaptive timing and temporal discrimination during associative learning. Neural Networks 2, 79–102.

Holder, M.D., Roberts, S., 1985. Comparison of timing and classical conditioning. J. Exp. Psychol.: Anim. Behav. Process. 11, 172–193.

Jenkins, H.M., Barnes, R.A., Barrera, F.J., 1981. Why autoshaping depends on trial spacing. In: Locurto, C.M., Terrace, H.S., Gibbon, J. (Eds.), Autoshaping and Conditioning Theory. Academic Press, New York, pp. 255–284.

Kamin, L.J., 1960. Acquisition of avoidance with a variable CS–US interval. Can. J. Psychol. 14, 1–6.

Killeen, P.R., Fetterman, J.G., 1988. A behavioral theory of timing. Psychol. Rev. 95, 274–285.

Killeen, P.R., Hanson, S.J., Osborne, S.R., 1978. Arousal: Its genesis and manifestation as response rate. Psychol. Rev. 85, 571–581.

Klopf, A.H., 1988. A neuronal model of classical conditioning. Psychobiology 16, 85–125.

Krebs, J.R., Biebach, H., 1989. Time–place learning by Garden Warblers (*Sylvia boria*): Route or map? Ethology 83, 248–256.

Libby, M.E., Church, R.M., 1975. Fear gradients as a function of the temporal interval between signal and aversive event in the rat. J. Comp. Physiol. Psychol. 88, 911–916.

Machado, A., 1997. Learning the temporal dynamics of behavior. Psychol. Rev. 104, 241–265.

Moore, J.W., Choi, J.-S., 1997. The TD model of classical conditioning: Response topography and brain implementation. In: Donahoe, J.W., Packard Dorsel, V. (Eds.), Neural-Networks Models of Cognition. Elsevier/North-Holland, Amsterdam, pp. 387–405.

Moore, J.W., Desmond, J.E., 1992. A cerebellar neural network implementation of a temporally adaptive conditioned response. In: Gormezano, I., Wasserman, E.A. (Eds.), Learning and Memory: The Behavioral and Biological Substrates. Lawrence Erlbaum Associates, Hillsdale, NJ, pp. 347–368.

Moore, J.W., Desmond, J.E., Berthier, N.E., Blazis, D.E.J., Sutton, R.S., Barto, A.G., 1986. Simulation of the classically conditioned nictitating membrane response by a neuron-like adaptive element: Response topography, neuronal firing, and interstimulus intervals. Behav. Brain Res. 21, 143–154.

Pavlov, I.P., 1927. Conditioned Reflexes (G.V. Anrep, Trans.). Oxford University Press, London.

Rescorla, R.A., 1967. Pavlovian conditioning and its proper control procedures. Psychol. Rev. 74, 71–80.

Rescorla, R.A., Wagner, A.R., 1972. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In: Black, A.H., Prokasy, W.F. (Eds.), Classical Conditioning II: Current Research and Theory. Appleton-Century-Crofts, New York, pp. 64–99.

Sakside, L.M., Wilkie, D.M., 1994. Time-of-day discrimination by pigeons, *Columbia livia*. Anim. Learn. Behav. 22, 143–154.

Schneider, B.A., 1969. A two-state analysis of fixed-interval responding in the pigeon. J. Exp. Anal. Behav. 12, 677–687.

Skinner, B.F., 1950. Are theories of learning necessary? Psychol. Rev. 57, 193–216.

Sutton, R.S., Barto, A.G., 1981. Toward a modern theory of adaptive networks: Expectation and prediction. Psychol. Rev. 88, 135–170.

Sutton, R.S., Barto, A.G., 1990. Time derivative models of Pavlovian reinforcement. In: Gabriel, M.R., Moore, J.W. (Eds.), Learning and Computational Neuroscience Foundations of Adaptive Networks. The MIT Press, Cambridge, pp. 497–537.

Tatham, T.A., Zurn, K.R., 1989. The MED-PC experimental apparatus programming system. Behav. Res. Meth. Instr. Comp. 21, 294–302.

Tesauro, G., 1986. Simple neural models of classical conditioning. Biol. Cybernet. 55, 187–200.